Control of gaze while walking: Task structure, reward, and uncertainty

Matthew H. Tong

Center for Perceptual Systems,
University of Texas at Austin, Austin, TX, USA

Oran Zohar

Center for Perceptual Systems,
University of Texas at Austin, Austin, TX, USA

Mary M. Hayhoe

Center for Perceptual Systems,
University of Texas at Austin, Austin, TX, USA

While it is universally acknowledged that both bottom up and top down factors contribute to allocation of gaze, we currently have limited understanding of how top-down factors determine gaze choices in the context of ongoing natural behavior. One purely top-down model by Sprague, Ballard, and Robinson (2007) suggests that natural behaviors can be understood in terms of simple component behaviors, or modules, that are executed according to their reward value, with gaze targets chosen in order to reduce uncertainty about the particular world state needed to execute those behaviors. We explore the plausibility of the central claims of this approach in the context of a task where subjects walk through a virtual environment performing interceptions, avoidance, and path following. Many aspects of both walking direction choices and gaze allocation are consistent with this approach. Subjects use gaze to reduce uncertainty for task-relevant information that is used to inform action choices. Notably the addition of motion to peripheral objects did not affect fixations when the objects were irrelevant to the task, suggesting that stimulus saliency was not a major factor in gaze allocation. The modular approach of independent component behaviors is consistent with the main aspects of performance, but there were a number of deviations suggesting that modules interact. Thus the model forms a useful, but incomplete, starting point for understanding top-down factors in active behavior.

Introduction

A fundamental problem in human vision is to understand the mechanisms that control attentional shifts from one location to another as visually guided behavior evolves in time. Many attentional shifts can be observed in overt shifts of gaze (Deubel & Schneider, 1996; Kowler, Anderson, Dosher, & Blaser, 1995), and it is this overt aspect of attentional shifts that we address here. A large body of work has focused on the way that both attention and gaze can be captured by novel stimuli, in particular on the nature of the stimuli that capture attention and gaze, and whether or not such capture is obligatory (e.g., Gibson, Folk, Teeuwes, & Kingstone, 2008; Jovancevic, Sullivan, & Hayhoe, 2006; Lin, Franconeri, & Enns, 2008). However, only part of the problem can be explained by such exogenously driven mechanisms. Human vision is in a large part goal driven. For example, in the context of even simple behaviors, such as walking, humans must accomplish a variety of specific goals, such as controlling direction, avoiding obstacles, and taking note of their surroundings. These competing demands for vision must be managed by selecting the necessary information from the environment at the appropriate time. However, the mechanisms that control goal-driven gaze shifts in the context of ongoing behavior are not well understood. What kind of a control structure is robust in the face of the varying nature of the visual world, allowing us to achieve these critical goals?

Although the neural circuitry underlying gaze shifts once a target is chosen has been intensively studied, we do not have much understanding of the control mechanisms that specify what location should be chosen as a target in the first place (Gottlieb, 2012). While it has long been recognized that behavioral goals of the observer play a central role in target selection (e.g., Buswell, 1935; Kowler, 1990), obtaining a detailed understanding of exactly how gaze targets are chosen
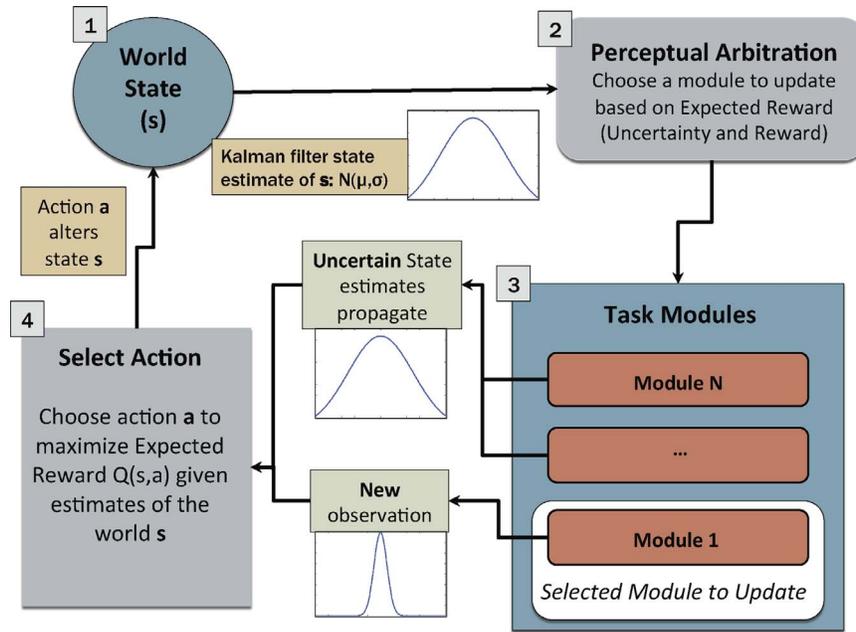
Figure 1. Flow diagram of the task-module architecture. Each task module keeps an estimate of its task-relevant state. Actions are chosen on the basis of the sum of reward that would be gained by all the modules. Better state estimates lead to better action choices. In the absence of direct observation, state uncertainty grows with time. The system uses gaze to update the state for the module that has the highest expected cost due to uncertainty. Kalman filters propagate state information in the absence of gaze updates (adapted from Sullivan, 2012).

on the basis of cognitive state has proved very difficult. Attempts to formalize this problem have typically taken the approach of explaining the effect of top-down factors by weighting a feature-based saliency map. For example, several models weight the stimulus saliency computations by factors that reflect likely gaze locations, such as sidewalks or horizontal surfaces, or introduce a specific task such as searching for a particular object (Kanan, Tong, Zhang, & Cottrell, 2009; Oliva & Torralba, 2006; Torralba, Oliva, Castelhano, & Henderson, 2006). Other models base top-down guidance on learned associations between features observed in locations that humans fixated when performing the tasks (Borji, Sihite, & Itti, 2011; Itti & Baldi, 2006; L. Zhang, Tong, Marks, Shan, & Cottrell, 2008). These models reflect the consensus that saccadic target selection is determined by activity in a neural priority map of some kind in areas such as the lateral intraparietal cortex and frontal eye fields (Bichot & Schall, 1999; Bisley & Goldberg, 2010; Findlay & Walker, 1999). However, the critical limitation of this kind of modeling is that it applies to situations where the subject inspects an image on a computer monitor, and this situation does not make the same demands on vision that are made in the context of active behavior, where visual information is used to inform ongoing actions. A broad range of different natural tasks have been investigated in the last two decades, and it is very clear that gaze is tightly linked, in time and location, to

the momentary task requirements, and often task demands can explain all but a few percent of the fixations (see Tatler, Hayhoe, Ballard, & Land, 2011, for a review). In attempting to understand these top-down effects, a critical limitation is that there is no formal representation of the task being performed. Whereas there have been successful attempts to model specific behaviors such as reading or visual search (Najemnik & Geisler, 2008; Reichle, Rayner, & Pollatsek, 2003), we need to develop a general understanding of how the priority map actively transitions from one target to the next as behavior evolves in time. The problem that we address here is how to capture the underlying principles that control these gaze transitions.

The challenge of modeling tasks is a difficult one, given the diversity and complexity of visually guided behavior. In this paper we consider a model introduced by Sprague, Ballard and Robinson (2007) that provides a general theoretical context for understanding the way that cognitive goals can influence gaze. We will examine the basic assumptions of this approach, and attempt to evaluate the empirical support for models of this kind. The structure of the model is illustrated in Figure 1. To simplify the problem, the model makes the assumption that complex behavior is composed of independent subtasks, or modules, such as avoiding obstacles, approaching objects, heading to a goal, and so on, and that specific information is gathered from the visual

image in order to perform the actions required for those tasks. Thus, information about obstacle location relative to the observer is required for an avoidance action, for example. Each subtask has an associated reward value that reflects the importance of the behavior for the agent, and allows the agent to learn how to arbitrate between the competing tasks to maximize the expected rewards learned using Reinforcement Learning (Sutton & Barto, 1998). At a given moment the subject acquires a particular piece of information for a module, using gaze (e.g., locates the nearest obstacle), takes an action (chooses avoidance path), and then decides what module should get gaze next. When a particular module is updated with information from gaze, as shown for Module 1 in the Figure, the new sensory information reduces uncertainty about the state of the environment relevant to that module (e.g., location of an obstacle). The next action is chosen on the basis of the learnt reward value associated with that action when in a given state by summing the rewards for all modules (Box 4). If a module's state is not updated, it is assumed that uncertainty about that state grows. For example, if an agent has just looked at an obstacle and updated its location, uncertainty about state relevant to another task such as location with respect to the goal is estimated from previous fixations, and uncertainty about that state grows with time. As a consequence of the action (e.g., moving in a particular direction), the state of the world is changed (Box 1), and the agent must decide which module's state should be updated next by gaze (Box 2). One hypothesis we will examine here is that the module given control of gaze is chosen on the basis of both the reward associated with the task (that is the subjective value of accomplishing the task) and the current perceptual uncertainty about the information needed to accomplish the task. In this formalization, task-driven vision is serial process, where one visual task accesses new information at each time step, and all other tasks rely on noisy memory estimates.

The model just described, which we will refer to as the Modular RL model, makes several central assumptions that seem like good candidates for explaining human performance. First is the assumption that complex behavior can be broken down into semi-independent subtasks, or modules. This makes the problem computationally tractable by reducing the size of the state space (Rothkopf & Ballard, 2010; Sprague et al., 2007). It also seems like a plausible approach to natural behavior, which seems well described as a set of sequential, task-specific operations (Ballard & Hayhoe, 2009; Hayhoe & Ballard, 2014; Land & Tatler, 2009). For example, when making a sandwich, successive fixations are tightly locked to the task in a way that makes almost every fixation interpretable in terms of

the current goal—e.g., fixate the knife handle while reaching to pick it up, fixate the peanut butter in the jar while scooping it out, fixate the tip of the knife while spreading, and so on. This tight linkage is discussed in Hayhoe, Shrivastava, Mruczek, and Pelz (2003), Land and Hayhoe (2001), and Land, Mennie, and Rusted (1999). The assumption that attention is deployed sequentially to different modules or subtasks is also consistent with the known limitations of attention, including the existence of a central attentional bottleneck that limits simultaneous performance of multiple tasks (e.g., Pashler, 1998), and experiments showing highly selective acquisition of information during a fixation (Droll, Hayhoe, Triesch, & Sullivan, 2005; Droll & Hayhoe, 2007; Rothkopf, Ballard, & Hayhoe, 2007; Triesch, Ballard, Hayhoe, & Sullivan, 2003). While this assumption is likely to be an oversimplification, it is a useful first step, as it reduces the gaze control problem to one of choosing which subtask should be attended at any moment. We test support for this assumption by manipulating the task structure subjects are given, making individual tasks or modules relevant or irrelevant. In the simplest case, individual modules should be active or inactive independently, with no interaction between modules.

Another central assumption of the Modular RL model is that tasks are prioritized on the basis of subjective value, or reward, and that this reward affects the eye-movements in addition to the choice of behavior. This assumption has strong support, given that neurons at many levels of the saccadic eye movement circuitry are sensitive to reward (see, for example, Gottlieb, 2012, for a review). In particular LIP neurons that are likely involved in saccade target selection have been implicated in coding the relative subjective value of potential targets (Trommershauser, Glimcher, & Gergenfurtner, 2011). There is also good evidence that the neural reward machinery acts in ways predicted by Reinforcement Learning models (Lee, Seo, & Jung, 2012; Schultz, 2000). However, reward effects in neurons have been observed with very simple choice response paradigms where the animal gets rewarded for looking at a particular target, whereas in natural vision, individual eye movements are not directly rewarded, but instead are for getting information that allows behavioral goals to be achieved, presumably with associated rewards. Thus, it is important to make the definitive link between the primary rewards used in experimental paradigms and the secondary rewards that operate in natural behavior. In human behavior, it has been shown that saccadic targeting is sensitive to explicit reward (money or points) in simple experiments involving a choice between a small number of targets (Navalpakkam, Koch, Rangel, & Perona, 2010; Schutz, Trommershauser, & Gegenfurtner, 2012; Stritzke, Trommer-

shauser, & Gegenfurtner, 2009). Targeting is also sensitive to implicit reward. For example, Jovancevic and Hayhoe (2009) demonstrated that while walking in a natural environment, subjects looked more frequently at potentially hazardous pedestrians who sometimes veered briefly toward them, than at pedestrians who simply stopped and so were visually salient but not hazardous. Since the events were of comparable visual salience, and the eye movements were anticipatory, before the pedestrian actually veered or stopped, this might be interpreted as reflecting the behavioral relevance or intrinsic reward value of the information. We will test this assumption by making various tasks more or less rewarding and look for the effects of this manipulation on gaze behavior.

Reward alone is not sufficient to understand why something becomes a gaze target. In the Modular RL model, uncertainty about task-relevant state of the scene is also important. In humans, there are various sources of uncertainty. For example, the lower resolution of the peripheral retina introduces uncertainty in the evaluation of sensory evidence, and indeed is the reason for fixating a target in many cases. If a subject has fixated a target and subsequently shifts attention, the information will be held in working memory, which decays over time. Further uncertainty will be introduced as the observer moves relative to the environment, since, in the absence of overt attention, locations will need to be updated by an estimate of the change in viewer location in the environment, and these estimates are subject to motor noise. The need to include uncertainty in the model to explain gaze choices stems from the fact that the optimal action choice is unclear if the state is uncertain. The Sprague model posits that the module or task that gets updated by gaze is the one that reduces reward-weighted uncertainty the most. In the original model, the world was static and uncertainty grew purely due to intrinsic factors like memory decay and motor noise; here we expand the experiment to include a dynamic and uncertain world.

There is mixed evidence for the role of uncertainty reduction in the choice of gaze target. Najemnik and Geisler (2005, 2008) showed that in the context of visual search for a simple pattern in noise, fixations appear to be chosen in order to reduce uncertainty. Similarly, fixations were governed by entropy reduction in a shape discrimination task (Renninger, Verghese, & Coughlan, 2007). Subsequent work by Verghese and colleagues, however, found that observers do not select targets to minimize uncertainty in a search for multiple targets, especially when under time pressure (Ghahghaei & Verghese, 2015; Verghese, 2012). We manipulate uncertainty by having some of our objects move in some conditions. If the model is correct, this should have an effect on subjects' gaze and walking behavior, showing an increased need for updated information

with increased looks to relevant objects and evidence that uncertainty grows since the object was last fixated.

The goal of this paper is to explore the validity of the approach outlined in the Sprague et al. (2007) model, in the context of a walking task. We developed a virtual environment where subjects walk along a path while avoiding obstacles and collecting targets. A similar environment was developed by Rothkopf et al. (2007) and demonstrated the importance of task relevance in gaze allocation. They compared their results with saliency predictions and random allocation and found that gaze allocation reflects primarily task relevance. Like that study, we manipulate task instructions as a proxy for manipulating reward, but we add different uncertainty conditions by adding motion to the objects. Another similar experiment, in a virtual driving environment, was performed by Sullivan, Johnson, Rothkopf, Ballard, and Hayhoe (2012). In that study participants were instructed to follow a lead car at a specific distance and to drive at a specific speed. Implicit reward was varied by instructing participants to emphasize one task over the other, and uncertainty was varied by adding uniform noise to the car's velocity. Gaze measures showed that drivers more closely monitor the speedometer when there is added uncertainty, but only if it was also associated with high task priority or implicit reward. Uncertainty reduction did not appear to affect task performance, however, as would be expected if uncertainty reduction allows better action choices. In the present experiment we reexamined this question. We also extend the examination of reward and uncertainty to a different domain from the Sullivan et al. (2012) experiment, by using a walking environment.

Our goal was to examine the basic assumptions of the model described above, to see if they are consistent with observed behavior; namely, is human gaze behavior affected by the uncertainty and reward structure as the model indicates? Our environment extended that of Rothkopf et al. (2007) by adding uncertainty to the obstacles, and by making the path more complex. The environment was designed to have observers meet similar visual challenges to those they might meet in the real world. To do this we had subjects follow a curved path across the room, avoid obstacles, and intercept targets. Since heading towards a goal, controlling the path walked, and avoiding and intercepting objects are all common demands of acting in the natural world, the goal was to make the experimental context as general as possible with respect to the class of behaviors involved in walking. We manipulate the relevance of the collection and avoidance tasks, by instructions, and observe the effects of this on gaze target selection. In particular we ask whether fixations to targets and obstacles increase in response to the task instruction (implicit reward), and
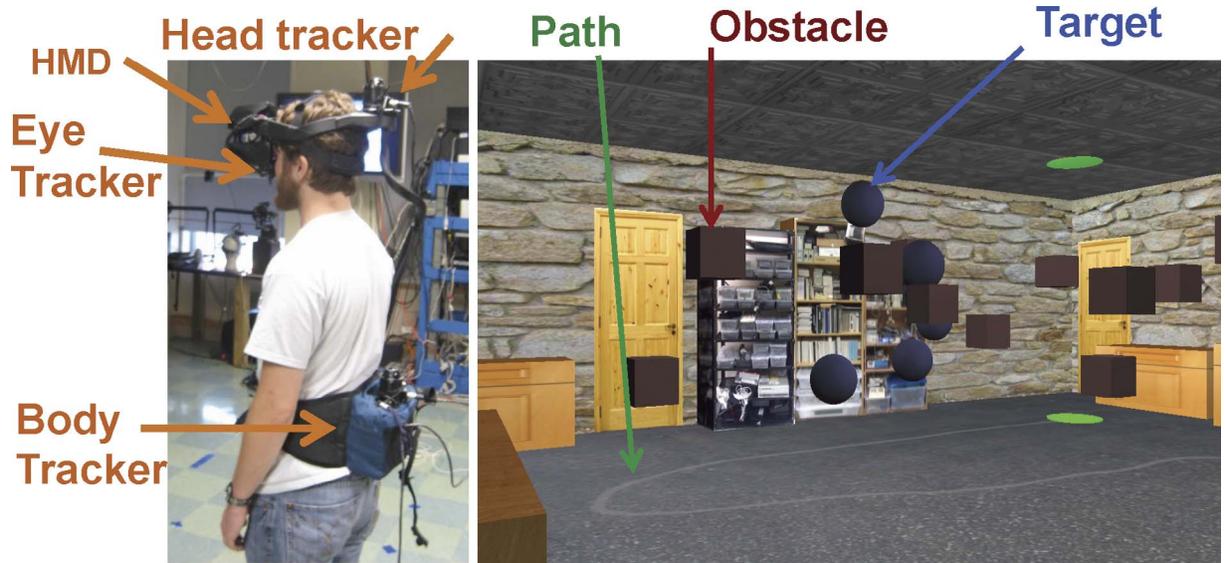
Figure 2. (left) The NVIS head mounted display, with Hi-Ball tracking system on the head and around the waist. (right) The virtual environment seen by the subject, showing the light gray path, the blue sphere targets, and the red/brown cube obstacles. The light green circle shows the location of the virtual elevator.

whether this increase in turn improves obstacle avoidance and target interceptions. We also ask whether increasing the positional uncertainty of the objects increases the frequency of fixations on task-relevant objects as suggested by the Modular RL model. We also test an alternative hypothesis that subjects simply look at the closest object. By combining different tasks, we examine whether fixations and action decisions are affected independently by the different task instructions as expected on the basis of the independent module assumption, or whether instructions on one task (say, avoid obstacles) influences fixations on irrelevant objects (e.g., increased target fixations). Although we find the expected effects of reward and uncertainty, we find deviations from complete independence of the modules. In addition we examine another assumption of the model: that subjects perform one task at a time while letting uncertainty grow in the other task.

Methods

Environment

To examine these questions we developed a virtual walking environment, illustrated in Figure 2. Subjects walked along or near the grey path shown in the Figure, while avoiding the red cubes and collecting the blue spheres that floated around eye height. This is similar to the paradigm used by Rothkopf et al. (2007), except that their path was wide and straight and the

objects were anchored on the ground. Having the objects float allowed us to introduce uncertainty by varying their position and provided greater distance between path and object information. The dimensions of the virtual environment were matched to the real room and subjects walked diagonally across the room, following the roughly S-shaped path. At each end of the path, a virtual elevator transported the subject to a new room with a different array of objects, and the subject then walked back in the opposite direction across the room. This allowed us to parse the experiment into a sequence of separate trials. There were 12 spheres and 12 cubes distributed randomly and uniformly around the room at heights between 4 ft. and 6 ft. For both targets and obstacles, a collision occurred if subjects passed within 6 in. of an object.

Apparatus

The virtual environment was created using Vizard software and was viewed using a NVIS SX-111 head mounted display. Each eye provides 76° horizontal and 64° vertical FOV at 1280 × 1024 resolution with a 50° region of horizontal overlap. The HiBall system by 3rdTech provided rotation and translation tracking at approximately 600 Hz for both head and body; one HiBall sensor was mounted to the back of the HMD while a second was attached to a pack subjects wore at their waist. The latency of the Hi-Ball signal is a few ms with high positional and angular accuracy. The whole system latency between a head movement and the screen being updated was 50–75 ms. The head mounted
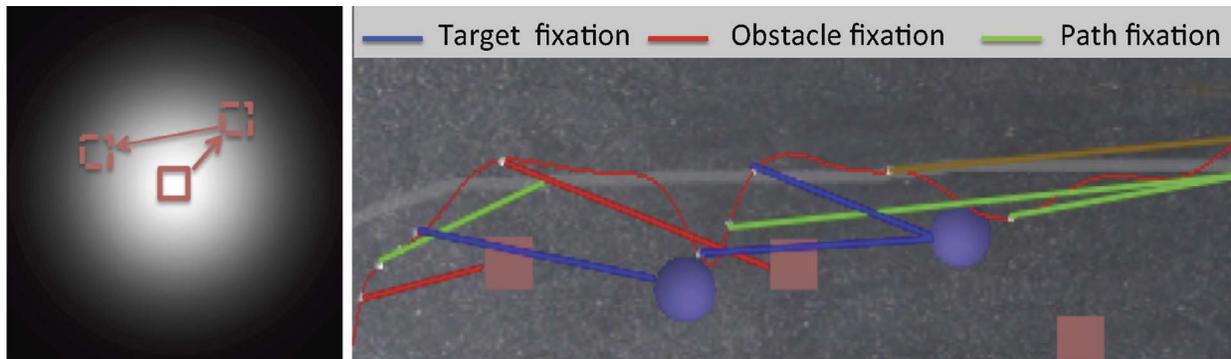
Figure 3. (left) An illustration of how objects moved in the high uncertainty condition. Objects were given a starting location in the room. New destinations were selected at random from a Gaussian distribution centered at that starting location with a horizontal standard deviation of 2′ and a vertical standard deviation of 1′. Objects then moved towards that destination at 1′/s. Upon reaching the destination, a new destination was sampled from the distribution, still centered at the initial location. Objects were thus constantly in random motion. In the low uncertainty condition, objects were stationary. (right) Top down view of a segment of the path showing the subjects trajectory (thin red line) and the fixations made (red for obstacle, blue for target, green for path, brown for background objects).

display was equipped with an Arrington Research View-Point eyetracker tracking the left eye at 60 Hz. Calibration was performed using a 9 point calibration grid at the start of the experiment. Calibration accuracy was checked again after the practice session and at the end of each block of trials, and the system was recalibrated when necessary. A video record of the scene camera and left eye camera was recorded at 60 Hz and then combined with metadata about head position, body position, and experimental conditions such as the location of objects in the room to generate a QuickTime video with attached metadata for subsequent analysis.

Procedure

As a "Reward" manipulation, we varied the task priority by varying the instructions. In different conditions, subjects were either instructed to just follow the path (Follow), or in addition to collect the blue target spheres (Collect), to avoid the red obstacle cubes (Avoid), or to both collect the targets and avoid the obstacles (Both). Thus, path following was included in all four conditions. Each floating object (target or obstacle) was therefore either task relevant or task irrelevant. Contact was made with an object by physically walking into it. When a target was contacted in the relevant condition, a fanfare sound was heard. When an obstacle was contacted in the relevant condition, a buzzer was heard. We assume that the fanfare sound is more pleasant than the buzzer. When either object was contacted in the irrelevant condition, a soft bubble pop was heard. The color of the targets and obstacles was counterbalanced in another version of the experiment in this environment (Hitzel, Tong,

Schutz, & Hayhoe, 2015) and was found not to affect the distribution of fixations. A similar control was performed in Rothkopf et al. (2007), so we did not repeat the control in the present experiment. In order to manipulate "Uncertainty," random motion was added to the objects in different conditions. In the High Uncertainty condition, objects moved along a straight path at 1′/s to a new location chosen from a three-dimensional Gaussian distribution centered at the object's original location and within the original range of heights. This Gaussian had a horizontal standard deviation of 2′ and a vertical standard deviation of 1′. When the object reached its destination, a new target location was chosen, and the object was continuously in motion, as shown in Figure 3 (left). In the Low Uncertainty condition the objects were stationary. Thus in all, there were four task conditions (Follow, Follow + Avoid, Follow + Collect, and Follow + Both) and four uncertainty conditions (high and low for each object type), making 16 conditions in total. Examples of a subject performing the Follow + Both task in both the Low Uncertainty and High Uncertainty conditions are shown in Video 1 and Video 2. The crosshair in the Videos shows eye position.

Subjects were given two trials of practice: one with moving targets and both objects relevant and one with moving obstacles and neither object relevant. They then performed 32 trials in blocks of eight. Each block had one of the four uncertainty conditions and had two consecutive trials for each task condition. The order of the tasks was Follow, Avoid, Collect, and Both. This order was chosen so as not to influence the single task conditions by doing the double task. Thus, it is possible there are some order effects. In another experiment in the environment (Hitzel et al., 2015) the order of the conditions was counterbalanced and no obvious order

effects were observed. The order of the uncertainty conditions was counterbalanced across subjects. At the start of each trial, subjects were verbally given instructions for the next room.

Participants

Subjects had normal or corrected-to-normal vision (contact lenses only). Subjects were UT undergraduates participating in the experiment for course credit and were naïve to the specific hypotheses being tested. Of 20 subjects run in the experiment, 12 completed the experiment and provided stable eye tracking data. The experiment was approved by the University of Texas Institutional Review Board, and all participants gave written informed consent.

Data analysis

Subjects' eye position data were analyzed using an automated system developed in-house. The eye signal was preprocessed using a median filter and then a moving average, both over three frames, to smooth the signal. Fixations were identified using a velocity criterion of less than 60°/s for at least 100 ms. (This is a relatively high threshold as the subject is walking, and the vestibular ocular reflex adds a slow component to eye-in-head velocity; compared with the velocity of a saccade, however, the compensation for the motion of head and body remains below this threshold, and this value has been found to reliably discriminate between saccades and steady gaze in previous experiments where Subjects walked in virtual environments (Jovancevic et al., 2006; Kit, Katz, Sullivan, Snyder, & Ballard, 2014; Li, Aivar, Kit, Tong, & Hayhoe, 2016; Rothkopf et al., 2007). If eye position remained within a 1.5° radius, transient velocity peaks of less than 80 ms were ignored and the data treated as belonging to a single fixation. By using the stored metadata and reconstructing the environment, we were able to recover the identity of fixated objects by intersecting the vector representing gaze direction with the 3D spatial volume of the object. To identify the location of the fixation in the scene, a 60 × 60 pixel window subtending approximately ±2° of visual angle, was centered around the location of the point of gaze on each frame, and each pixel in the window returned a label for the type of object it contained. Scene regions were labeled as one of the following: Path, Obstacle, Target, Elevator, or Background. As we tracked only the left eye, we are unable to label gaze to points outside the field of view of the left eye—such frames were labeled as Background. Only a small number of fixations fall in the corresponding monocular region of the left eye, so by

inference most fixations are within the binocular field of view. The fixation target on each frame was labeled as the one with the greatest proportion of pixels, except when the background was the greatest proportion. In this case, the next most frequent object was chosen. (This avoids failure to identify object fixations when gaze was near an object's edge with much of the fixation window off the object.) The region labeled as the fixation target for most frames during a fixation was labeled the target for whole fixation. Consecutive fixations to the same nonbackground target were combined if the time between them was less than 80 ms to further reduce noise. The Path region was defined as a 2′ strip around the actual path, which was a thin line in the scene. The automatically identified labeled fixations were validated against a frame-by-frame manual analysis of the video records of a subset of the data. The manual coding allowed fixation analysis parameters to be optimized to best match the manual labeling, and the manual and automated analyses were in good agreement. A similar technique for fixation analysis was used in Kit et al. (2014) and Li et al. (2016). In our analysis, fixations were measured in the period following movement initiation at the beginning of the path and prior to reaching the next elevator, when subjects were farther than half a meter from the start and end elevators.

Results

The goal of the experiment was to examine the influences on both eye and body movements as subjects walked through the environment. The focus was on the factors that control momentary choices of what is the most important action from moment to moment. Figure 3 (right) shows a segment of gaze and walking behavior that illustrates these momentary gaze and action choices. In this segment, the subject makes a fixation to an obstacle (red), followed by avoidance of the obstacle. During this avoidance path, the subject makes a path fixation (green), then a fixation to a target (blue), that the subject detours to collect. How are these target decisions made?

Gaze allocation: Task effects.

First we examine the effect of the task instruction on the distribution of fixations. Figure 4 plots the number of fixations on targets, obstacles, and path for the four task conditions. Subjects completed the trials fastest in the follow condition (19 s, with 31 fixations on average), increasing to about 24 s with 42 fixations for the Avoid or Collect tasks, and 28 s with 48 fixations
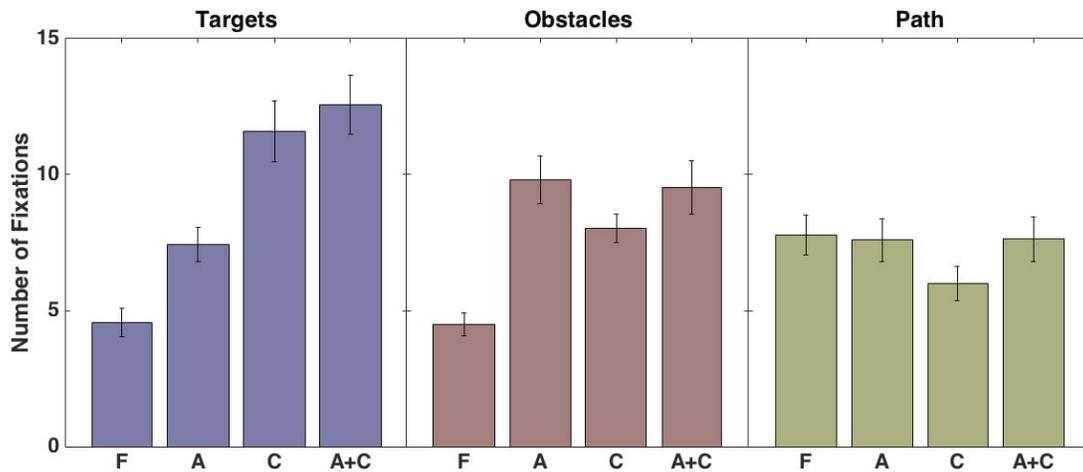
Figure 4. Task effects on number of fixations on targets, obstacles, and path. F indicates the Follow instruction; A, the Avoid and Follow instruction; C, the Collect targets and Follow instruction; and A+C is all three. Error bars indicate between-subjects standard errors.

for the full Avoid and Collect task. Note that in this and in subsequent plots we have chosen to plot fixation number rather than proportion, but none of the conclusions drawn are affected by this choice. We chose fixation counts as it seemed a more direct reflection of the effect of instructions and object motion on behavior. Subjects often fixated a target and maintained fixation until collection or shortly before; to avoid this behavior giving too much influence on our measure, we consider the fixation count rather than duration and combine multiple temporally adjacent fixations to the same object. Data are averaged over all uncertainty conditions for each subject. Error bars reflect between subjects variability. Background fixations (not landing on any object of interest) account for approximately 40% of fixations in all conditions and are not plotted. As expected, and as found previously in a similar environment (Rothkopf et al., 2007) the distribution of fixations is sensitive to the task goals. Target fixations are highest in the Collect and Avoid + Collect conditions, and obstacle fixations are highest in the Avoid and Avoid + Collect conditions. Fixations to targets when task-relevant are a little more frequent than to obstacles when relevant. This may result from the different behavior during collection versus avoidance. During collection subjects orient towards the target and maintain gaze on the target whereas when avoiding obstacles, subjects direct themselves away from obstacles. For this reason, we focus primarily on how task affects fixations for a given object rather than differences between fixation numbers on targets versus obstacles. Path fixations are relatively constant, at about seven fixations, except for the collect condition, perhaps as this involves greater deviations from the path with the targets largely dictating the trajectory through the space. While the task modulation is substantial, both targets and obstacles each get a

number of fixations, even in the Follow condition, where they are not explicitly relevant. If we interpret fixations as reflecting some intrinsic value for the information, the instructions are only one factor in determining what subjects consider relevant while walking. Presumably knowing the structure of the scene and the location of objects in the environment is something that has been rewarded during ordinary visual experience, as it is almost always important. Alternatively, these fixations might reflect more stimulus-driven saliency mechanisms, although we present evidence against this possibility below (see Figure 5). Finally, if an object intersected the line of sight either through subjects' movements or its own, it would likely be counted as a fixation, so some random noise may also be a contributing factor.

Note that target fixations go up when subjects are instructed to avoid obstacles, even though targets are not explicitly relevant. That is, the number of target fixations goes from less than 5 in the Follow condition to about 7.5 in the (Follow and) Avoid condition. Similarly, obstacle fixations go from less than 5 in the Follow condition to 8 in the (Follow and) Collect condition. Both these differences were highly significant on a matched pairs t test: $p = 0.003$, $t(11) = -3.74$ and $p = 0.0002$, $t(11) = -5.59$ respectively. This suggests that the presence of targets affects obstacle fixations and similarly that the presence of obstacles affects target fixations. In the simplest form for the Modular RL model, described by Sprague et al. (2007), the amount of reward associated with a module affects its fixations, and increasing the reward of one module should either not change the fixation behavior of irrelevant objects from their baseline or decrease the number of fixations due to competition for a limited resource. This increase of gaze to irrelevant objects therefore suggests that
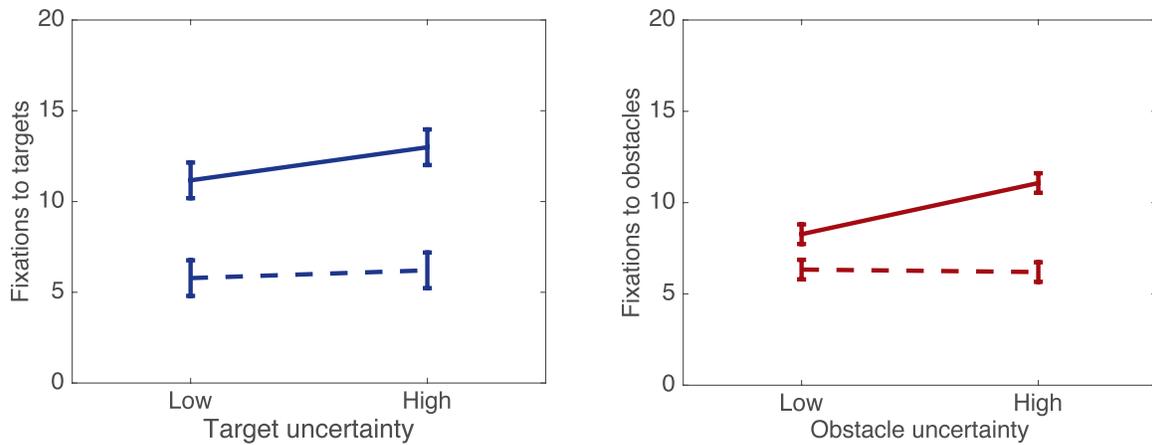
Figure 5. The effects of uncertainty and relevance on fixations. The left side of the figure shows number of fixations to targets for low and high target uncertainty. The right side shows number of fixations to obstacles for low and high obstacle uncertainty. The solid line shows the condition where targets (left) or obstacles (right) are relevant, and the dashed lines show the corresponding irrelevant condition.

modules are not completely independent. We take this issue up in the Discussion.

Uncertainty manipulation: Effects on gaze

As discussed in the Introduction, the role of uncertainty on gaze allocation is unclear. While it seems clear that in general gaze is deployed to get necessary information, just how this functions in a dynamic behavioral context is not understood. Here we manipulated the amount of extrinsic noise by adding motion to the targets and obstacles. Does this lead to more frequent obstacle and target fixations? This is shown in Figure 5. Data were averaged for each subject for the High and Low Uncertainty conditions, and for both task Relevant and Irrelevant conditions, and the Figure shows between-subjects variability estimates. A four-way ANOVA of the subject identity, object identity, relevance, and uncertainty showed a main effect of an object's relevance, $F(1, 45) = 145.02$, $p < 0.001$; uncertainty, $F(1, 45) = 9.73$, $p = 0.0032$; subject identity, $F(11, 45) = 6.47$, $p < 0.001$; and object type, $F(1, 45) = 7.34$. $p < 0.001$. There were also significant interactions between relevance and uncertainty, $F(1, 45) = 7.56$, $p = 0.0086$; subject identity and relevance, $F(11, 45) = 5.67$, $p < 0.001$; and object type and relevance, $F(1, 45) = 11.54$, $p = 0.0014$. The main effects of relevance and uncertainty are as predicted by the model. We further examine the relevance and uncertainty interaction. When obstacles were relevant to the task, subjects increased the number of fixations to obstacles when obstacle uncertainty was high, and obstacles were relevant to the task: from 8 to 11, $t(11) = -2.692$, $p = 0.021$. The uncertainty manipulation has no significant effect when obstacles were not relevant:

approximately six fixations, $p = 0.813$, $t(11) = 0.242$. In the case of target fixations, the effect of adding positional uncertainty increased fixation by about 2 (from 11 to 13) when targets were relevant. This increase was significant, $p = 0.043$, $t(11) = -2.286$. Again, there was no effect of the uncertainty manipulation when the targets were not relevant, $p = 0.210$, $t(11) = -1.330$. It is of interest that for both targets and obstacles, the addition of positional uncertainty to irrelevant objects did not increase fixations to those objects. Thus, the added motion in peripheral vision did not attract gaze, as might be expected from most models of saliency where motion is a highly salient cue (Borji & Itti, 2015). Interestingly, there was a suggestion that the motion of the irrelevant object increased looks to relevant objects, although this did not achieve significance (ANOVA on the Collect and Avoid trials looking at the effects of the uncertainty of the irrelevant object, $p = 0.159$, $F(1, 45) = 2.05$.

Alternative strategies

One central premise of the Modular RL model is so far supported by the gaze data, namely, that fixations are at least partially determined by both task priority (reward) and uncertainty. However, it is possible the data might be consistent with alternative hypotheses. For example, in the simplest case, subjects might simply look at the closest object to them as they walk along the path, and depending on how subjects position themselves this might potentially yield the same results shown above. These two hypotheses overlap some because proximity has a direct effect on how uncertainty and reward interact; if a target is on the other side of the room, its exact position is largely irrelevant
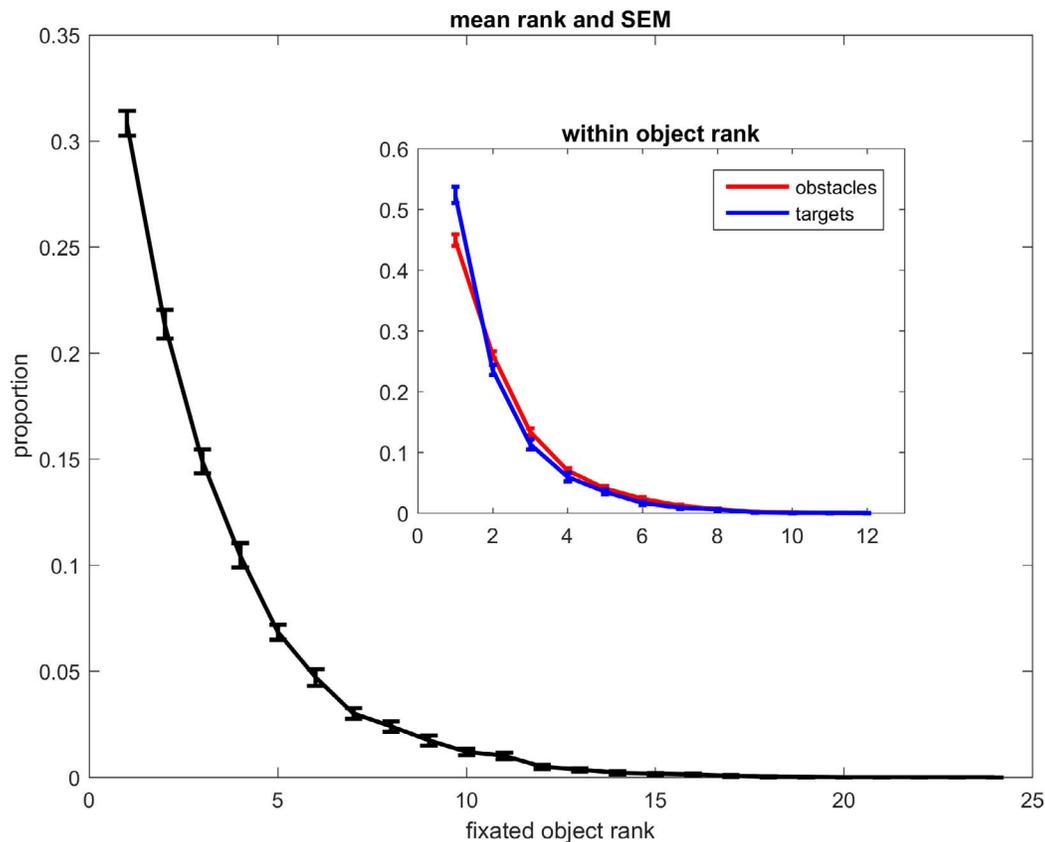
Figure 6. Proportion of fixations as a function of object distance from the subject, expressed as a rank. Thus 1 is the closest object on screen, 2 the next closest, and so on. Error bars are between-subjects standard errors. The inset shows the distributions separated by object type.

as one can head in the right general direction, whereas collecting a nearby target requires having an accurate representation of its position. To examine this competing hypothesis, we counted the frequency with which subjects looked at the closest object within the field of view, at the second closest, and so on. The resulting frequency distribution is shown in Figure 6. The figure shows that about 30% of the fixations are to the closest object, and substantial proportions of the fixations go to the second and third closest objects. Thus there is a clear tendency to look at nearby objects, but not simply at the closest one. Thus, the simple heuristic of looking at the closest objects does not match the data, supporting the hypothesis that task and uncertainty play a role.

The inset of the figure has similar histograms, but now separated out into object-specific rank. About half the time subjects look at the closest obstacle on screen and half the time at the closest target, presumably as dictated by the current attentional focus. This indicates that the simplifying assumption used by Sprague et al. (2007) that modules receive updates on the closest relevant object is approximate but not perfect, and that subjects may be planning further ahead. Thus the model should be extended to allow the tracking of

multiple objects, so that the uncertainty associated with the more distant objects could be factored in to gaze choices.

Another consideration is that fixation locations tend to be localized around the center of the field of view. This is most well known when viewing stimuli on a screen, where there are frequently strategic advantages to not focusing gaze at the edge of a display; photographers tend to place interesting content towards the center of an image and locations beyond the edge of the screen reveal nothing about the scene portrayed (Tatler, 2007). Eye movements in the real world do not have the edge of the screen driving gaze to the center, but nevertheless tend to show a central bias with the majority of fixations falling within the central 20°–30° of view relative to the head, but showing a large effect of task (Foulsham, Tong, & Kingstone, 2011; 't Hart & Einhäuser, 2012). Given that subjects approach targets for collection, this bias could mean that simply walking towards targets would increase the number of fixations to targets, driving the task effect we observe. However, as shown in Figure 3 (right), the subjects' fixations tend to precede the turning action towards a target or away from an obstacle. Furthermore, the opposite would be true for obstacles; avoiding obstacles
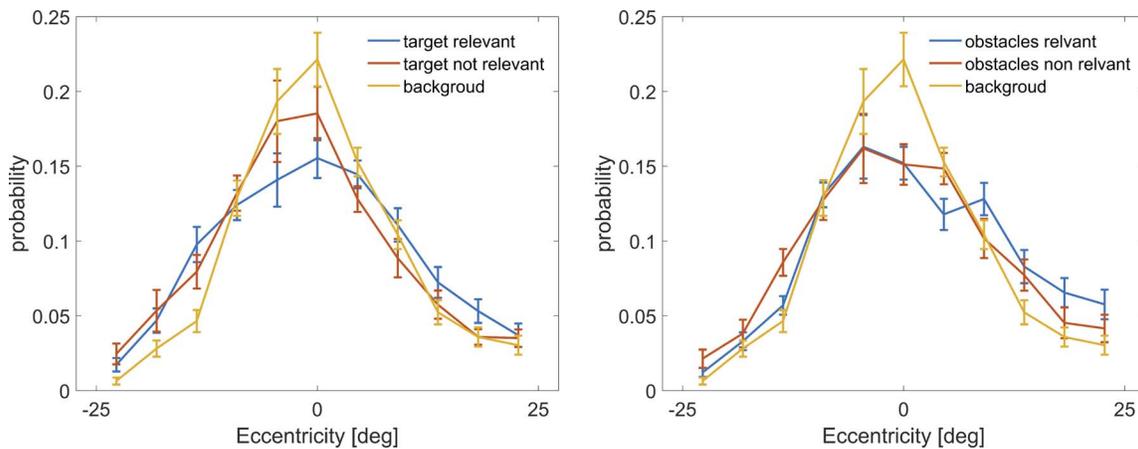
Figure 7. The distributions of fixations along the horizontal dimension for targets, obstacles, and background. (left) Target fixations in both relevant (blue) and irrelevant (red) conditions compared with background (yellow). (right) Obstacle fixations in both relevant (blue) and irrelevant (red) compared with background (yellow).

would lead to a reduction in the time an obstacle is straight ahead of a subject and a reduction of fixations due to the central bias; however, active avoidance increases the number of looks to obstacles. To examine this issue further, we plotted the distribution of fixations across the visual field to targets, obstacles, and to the background (defined as fixations not to the path, obstacles, or targets). Figure 7 shows the distribution of these fixations across the horizontal axis. While fixations exhibit a central bias, they are broadly distributed across the visual field with subjects making many large eccentric fixations. Background fixations are more narrowly distributed horizontally in the visual field compared with obstacle and target fixations. We also observe an effect of task for targets, where the distribution of targets when being collected is wider than when they are not task relevant; however, there does not seem to be a significant effect of relevance on

the fixation distribution for obstacles. All of these distributions were significantly different according to a Kolmogorov-Smirnov test ($p < 0.01$). The effects observed go beyond what is able to be accounted for with the central bias and reveal an active search for objects of interest.

Action decisions

In our conceptualization of natural behavior, gaze is used to update internal estimates of world state in order to make decisions that satisfy momentary task goals. Thus, gaze choices should be related to action decisions. Figure 8 (left) illustrates a subject's path when walking through the environment in the different task conditions, and Figure 8 (right) shows collection and avoidance performance when doing the four tasks,
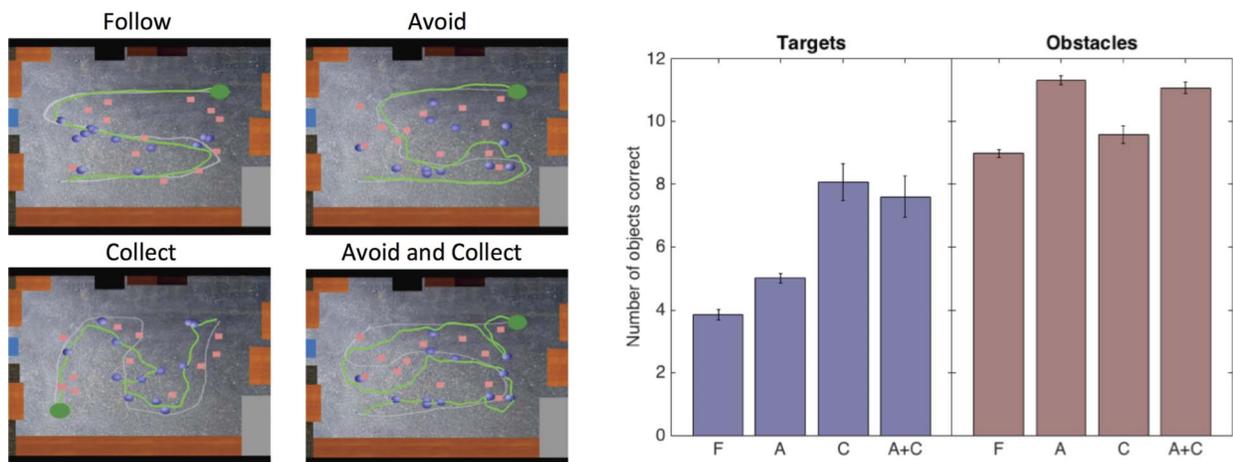

Figure 8. The effects of explicit task on performance. (left) Example paths taken by a subject under the different instructions. (right) The number of targets collected and obstacles avoided are plotted for the for different task instructions. F = Follow, A = Avoid, C = Collect, and A+C = Both. Error bars are standard errors between subjects.
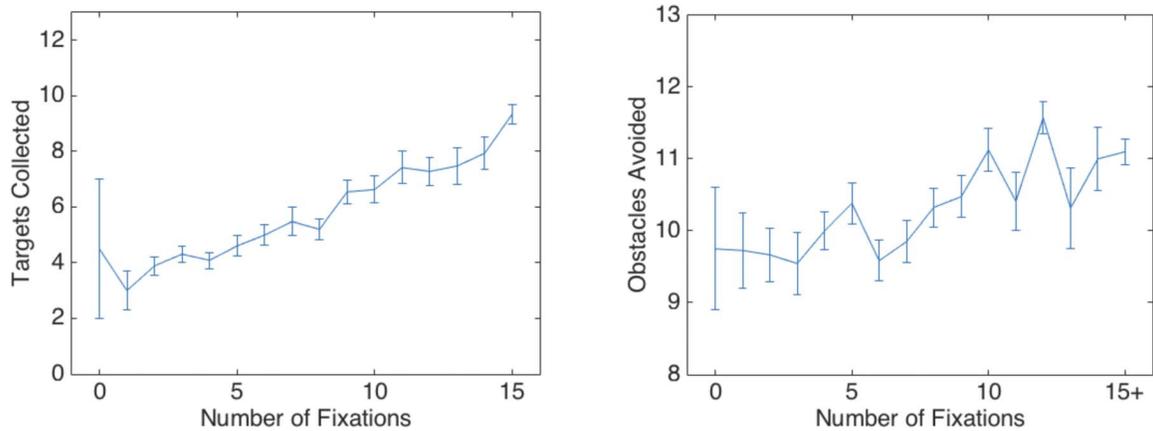
Figure 9. (left) Target collections as a function of number of fixations to targets. (right) Obstacles avoided as a function of number of fixations to obstacles. Data were averaged over all conditions for a given number of fixations, and error bars reflect the standard error of the estimate of the mean for that number of fixations.

averaged over uncertainty conditions. Subjects collected more targets when given the Collect instruction. Likewise, subjects avoided more obstacles when instructed to Avoid. Note that about four targets were collected, and nine of the 12 obstacles were avoided (that is, three collisions) even when simply instructed to follow the path. As targets and obstacles were distributed at random, the greater number of targets collected (four) than obstacle collisions (three) suggests that the label itself induces slightly different task-specific behavior, $p = 0.0090$, $t(11) = 3.16$. Additionally, target collection increased slightly in the Avoid task, $p = 0.0003$, $t(11) = -5.13$, and avoidance increased slightly in the Collect task, $p = 0.0002$, $t(11) = -5.6$. Thus, there is some value to collection and avoidance even when not explicitly instructed, and instructions to avoid influence the value of collection. An alternative interpretation is that the task execution is not strictly independent as postulated in the Modular RL model. Such a failure of independence could result from some kind of global scene analysis, for example, rather than independent valuation of targets and obstacles.

Relation between gaze and actions

Given that subjects' behavior is modulated as expected, we now address the issue of the relation between gaze and actions. Figure 9 plots average targets collected and obstacles avoided as a function of the proportion of fixations. The effect of fixations is clear in the case of target collection. The more often targets were fixated, the more often they were collected, regardless of task relevance (correlation coefficient of 0.63, $p < .0001$). The trend is less pronounced for the obstacles, since few obstacles are collided with. However, the more obstacles were fixated, the fewer

collisions with obstacles there were, and the correlation is significant (correlation coefficient of 0.22, $p < 0.001$). The correlation for targets is perhaps unsurprising, since subjects approach targets when collecting them making a fixation to them likely, but the opposite is true for obstacles so taken together the two correlations provide evidence for the coupling of gaze and behavior.

We also noted that other aspects of gaze behavior were specific to the particular task goal, as shown in Figure 10. Thus, the mean distance at which subjects started a fixation to the path was very stable, at about 3 m horizontal distance, whereas fixations to targets and obstacles were initiated when objects were at about 1.7 m and 1.8 m horizontal distance respectively. We performed a two-way ANOVA on the mean fixation distances for each subject for targets and obstacles in the high and low uncertainty conditions. The difference between the relevant and irrelevant conditions was significant, with fixations to obstacles being made when obstacles are closer in the relevant condition, $p < 0.0004$, $F(1, 44) = 14.78$, and for fixations to targets when targets are relevant, $p = 0.0061$, $F(1, 44) = 8.29$. This suggests that the distance at which objects were fixated reflects the behavioral needs of the subtask. There was also an effect of uncertainty on the distance at which targets and obstacles were viewed: two-way ANOVA with object and uncertainty, $p = 0.005$, $F(1, 44) = 8.7$, as shown in Figure 10 (right). In the high uncertainty condition, the average fixation distance is about 0.3 m closer for obstacles, $p = 0.004$, $t(11) = 3.64$, while the difference for targets was about 0.1 m and did not achieve significance, $p = 0.403$, $t(11) = 0.87$. The distance an object was last fixated is a good indicator of how long it will be until the subject passes the obstacle, and is correlated strongly with the time taken until the subject comes closest to the obstacle ($r = 0.74$ for low uncertainty, $r = 0.63$ for high uncertainty).
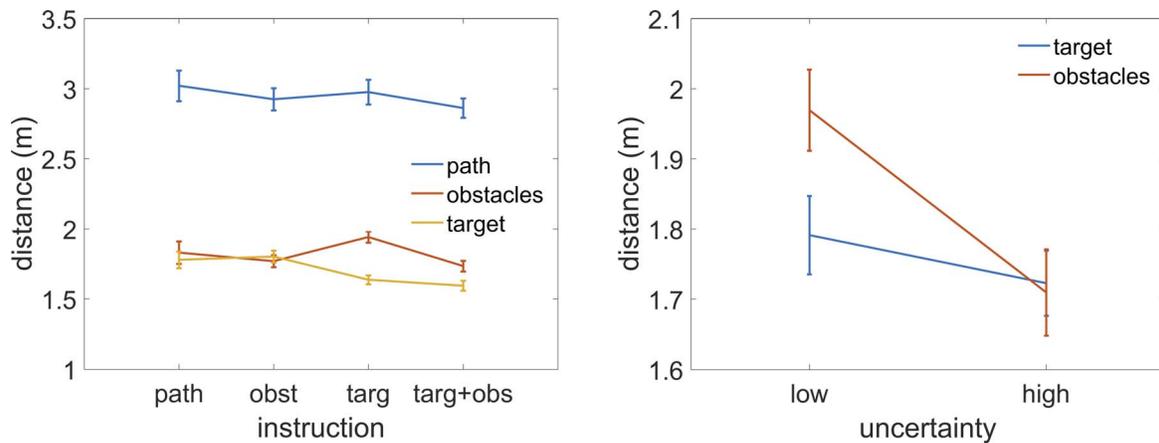
Figure 10. (left) Average horizontal distance of fixations on path, obstacles, and targets as a function of instruction. (right) Effect of Uncertainty on fixation horizontal distance for targets and obstacles. Error bars are standard errors between subjects. Distances were measured at the start of each fixation.

Discussion

We observed gaze and walking performance in an environment that presented similar challenges to those present in the everyday behavior of walking in a cluttered environment with moving objects. Our goal was to work towards an understanding of gaze control in complex behavior that integrates gaze into task performance. We examined the assumptions of the Sprague et al. (2007) Modular RL model, and extended the experimental domain from previous similar studies. Overall the data are consistent with the assumptions of the model, although there are a number of complexities not explained by this simple framework. We examine these assumptions in turn.

Task relevance and reward

Subjects must have some mechanism whereby behavioral goals determine where they attend, and thus where they look. We suggest here that in many natural contexts there are a number of goals that must compete for gaze, and that both gaze and walking direction choices reflect a single value for that goal relative to the competing goals. We attempted to manipulate this value using instructions, and find that gaze allocation and walking direction choices both reflect the effect of instructions, and by inference, implicit reward, as expected. A simple strategy of looking at the closest object of either kind was not able to account for the results. Thus, subjects execute searches to find the next target for gaze based on a choice of which task to attend to. We also showed a connection between eye movements and task performance. For both targets and obstacles, making more looks to a task relevant object was correlated with improved performance of

that task, meaning that more looks to targets correlated with more target collections and more looks to obstacles correlated with more obstacles being avoided.

It might be argued that instructions are a weak way to manipulate some underlying reward parameter. Indeed, it would be preferable if subjects could reveal stable intrinsic preferences for particular goals such as obstacle avoidance. It has been demonstrated that walking decisions in a similar experimental environment can in fact be used to infer an underlying value for a particular subtask, and these values appear to be quite similar across different subjects (Rothkopf, Hayhoe, & Ballard, submitted; Tong & Hayhoe, 2014). Interestingly, Rothkopf et al. (submitted) have shown that the values of the Collect and Avoid tasks estimated separately can account for the values of Collect and Avoid when performed together in the Both condition. This suggests that the instruction manipulation is effective in changing some fairly stable internal parameter reflecting subjective value. As discussed in the Introduction, there is a growing body of literature indicating the effect of both extrinsic and intrinsic reward on gaze behavior. What we add here is evidence that the relative value of different goals can be used to arbitrate between momentary choices to control both gaze target and walking direction. In the present conceptualization, gaze target is chosen to reduce uncertainty in order to make better decisions about walking direction, and in the Reinforcement Learning framework, walking direction is chosen to maximize future expected discounted reward that has been learned through experience. As pointed out by Eckstein, Schoonveld, Zhang, Mack, & Akbas (2015), and others (Gottlieb et al., 2012; Gottlieb, Hayhoe, Hikosaka, & Rangel, 2014) eye movements themselves are not directly rewarded in the context of behavior, unlike many of the neurophysiological experiments as

well as some of the psychophysical experiments (e.g., Navalpakkam et al., 2010, Schutz et al., 2012; Stritzke et al., 2009). Eckstein et al. (2015) argue that, at least in some cases, it is the decision informed by the eye movement that gets the reward, not the eye movement per se. In the present context the decisions are easily observable in the choice of walking direction, so we are able to see a more direct link between the gaze choices and the walking decisions made.

Uncertainty

By adding an external source of uncertainty in the current paradigm, we were able to increase the number of fixations on objects in the environment. This is true only when the object is task relevant. The effect of uncertainty is expected on the basis of the Modular RL model, although the effect was most apparent in the case of obstacle fixations. It is not entirely clear why the effect is stronger for obstacles, although it may indicate that the two tasks are different in ways not captured by the simple conceptualization in the model; for instance, since targets are approached, they are more likely to be fixated while they're being collected regardless of their uncertainty. The relatively modest effect of added uncertainty suggests that internal sources of uncertainty may dominate in the present circumstances, and these are not reflected in the figure. The added uncertainty had no significant effect on fixations to the moving object when the task was irrelevant. A similar finding was observed by Sullivan et al. (2012) in their driving paradigm, where uncertainty was added to the speedometer; however, there the added uncertainty was not visibly distinctive. By introducing uncertainty using motion, one might expect that irrelevant objects would be more salient in the high uncertainty condition, but we did not find this to be the case; it was only when objects were relevant to the task at hand that movement increased the number of fixations. In contrast, many bottom-up models predict high salience for unpredictably moving objects (Itti & Baldi, 2006; Peters & Itti, 2006) and models that tune attention to feature channels to provide top-down guidance would tend to keep motion and color cues separate (e.g. Navalpakkam & Itti, 2005), posing a challenge for how a moving relevant object would capture more attention but an irrelevant object would not.

Uncertainty reduction as a controlling factor in gaze target selection has also been investigated by Ghah-ghaei and Verghese (2015) and Verghese (2012). In these investigations, subjects performed a search task in noise for multiple targets, and it was observed that subjects choose the most likely location as often as the most uncertain location. It is not entirely clear how to compare these results with ours. One difference is that

we pose the question in a somewhat different manner, in that we are looking for more frequent object fixations when the locations of those objects is more variable, and subjects need the location information in order to perform an action, whereas in Verghese and colleagues' experiments the question is whether subjects look at a location that reduces uncertainty about whether a target is there or not. Another factor is that implicit reward may influence the behavior in complex ways and make interpretation more difficult. It may be that fixating targets may be more rewarding than resolving the uncertainty since those targets are being acted upon in the next step, selecting them and receiving a reward. Finally, Ghahghaei and Verghese (2015) noted that the effects of uncertainty become stronger as subjects have longer to make their decisions, and in our context subjects were not under time pressure. Thus, the specific experimental context and the nature of the behavioral goals may need to be taken into account when evaluating the importance of uncertainty in gaze target choice.

In our experiment, uncertainty effects are revealed not only in gaze choices, as in Verghese and colleagues' experiments, but also in the walking decisions subjects make. We found that as the number of fixations increased in the various conditions of the experiment, performance in target interceptions and obstacle avoidance improved. This link is important in validating a model like Modular RL, since it posits that the information acquired by gaze is used to inform action decisions. In Sullivan et al.'s (2012) driving experiment, measures of performance did not reflect the different distribution of fixations when uncertainty was added. We also observed that increasing uncertainty shortened the distance at which subjects fixated the objects. It is possible that this is caused by the greater fixation frequency after subjects first fixated an object, the necessary information being available later in time, or that subjects are adopting a strategy to reduce uncertainty about task relevant state information. All of these explanations are in agreement with the model, but it would be worthwhile to better understand the cause.

One important aspect of the current theoretical context is the idea that the state for only one module is updated at a time, and the value of other task-relevant variables is propagated in short-term memory, with noise accumulating as a function of time. Previous findings by Droll et al. (2007) in a block manipulation task, indicated that subjects do not necessarily update representations of the objects they are manipulating if they are attending to other aspects of the task, and instead use information stored in memory rather than current (foveal) sensory data. In the present study, task performance was correlated with fixations to task-relevant objects. The fact that the fixation distance on

objects decreases in the high uncertainty condition is consistent with the notion that uncertainty for unattended modules grows with time. Note, too, the illustration in Figure 3 (right), where the avoidance path occurs while the subject is fixating the next target for collection, suggesting that once a plan is underway and the information it needs determined, that attention can move on. Thus, the present experiment is consistent with previous work discussed above, showing that subjects allow for both intrinsic and extrinsic uncertainty when executing actions, but in addition it suggests that subjects typically take account of the fact that less recently acquired sensory data will have greater noise as the representation decays with time in visual memory and plan actions accordingly. Previous studies have also suggested that subjects take into account intrinsic uncertainty in visually guided action (Faisal & Wolpert, 2008; Körding & Wolpert, 2004; Maloney & Zhang, 2010; H. Zhang, Daw, & Maloney, 2015).

Modules

One of the simplifying assumptions of the Sprague et al. (2007) model is that complex behavior can be decomposed into independent subtasks or modules that access relevant parts of the state space independently, and that each fixation provides an update to a single module. These modules are also assumed to be able to access the information they need efficiently; when the visual system determines that the target module should receive a state update by making a fixation, a target will be fixated. This simplifies the modeling significantly and is motivated by the sequential, task-directed behavior one observes when performing real world tasks (Ballard & Hayhoe, 2009; Hayhoe & Ballard, 2005; Land & Tatler, 2009). Several of our findings provide challenges to this account.

First, we assume that objects are either relevant or irrelevant on a trial-by-trial basis following the instructions given, but there's some evidence that subjects did not strictly follow this (see Figure 8). Even when subjects were told to simply follow the path, they tended to make more contacts with targets than obstacles, suggesting some preference beyond the current instructions. More interestingly, when subjects were told to collect targets, they increased their avoidance of obstacles; being told to avoid obstacles likewise improves target collection. This means that the analysis based on objects being considered task irrelevant is not strictly correct; it appears that the floating objects here maintain some relevance for the subject even when irrelevant to the experimenter-specified task, and that their subjective relevance may vary. Inverse Reinforcement Learning can be used to recover the implicit reward structures subjects actually are using (Rothkopf et al., 2010), and some preliminary results show just those kinds of recovered weights (Tong & Hayhoe, 2014). In any event, relevance here could perhaps be better understood behaviorally as high versus low relevance conditions, not relevant versus irrelevant.

There also is some interaction between modules when one looks at eye movements. First, the instruction to avoid obstacles led to an increase in target fixations, and vice versa. There are several potential reasons for this interaction. This matches the basic pattern discussed above in the behavioral data, so it may be due to "irrelevant" objects still having some implicit reward, but the effect is much more pronounced, and one does not see the effect of increased uncertainty driving more looks towards irrelevant object, as one would expect if viewing such objects was rewarding. The presence of irrelevant objects could also make the visual search task more complex, as a consequence of crowding, since all floating objects were at approximately the same height. This crowding could have also made the proper labeling of fixations more confusable and noisy, although the automatic labeling was verified with hand coding. Alternatively, subjects might simply search for objects of any kind, and then move on if the irrelevant object is encountered, meaning the influence of task occurred after a fixation rather than when selecting them; however, the fact that more fixations were made to objects that were task relevant makes this conclusion unlikely. Lastly, while looks to the path appear constant across most tasks, during the Collect task, fewer path fixations are made. This might result if the targets are providing the main waypoints of the path, with subjects merely being concerned with keeping their trajectory vaguely close to the path. A similar interaction was observed by Rothkopf et al. (2007), who found that fixations on obstacles served the dual purpose of locating obstacles and the subject's position on the sidewalk, reducing the number of fixations to the sidewalk. Taken as a whole, the results support the modular approach in that adjusting the relevance and uncertainty of each task had significant effects on both eye and walking behavior. However, the naïve approach of assuming complete independence falls short of being able to capture all aspects of behavior.

Background fixations and other factors

It is important to note that a substantial proportion of fixations were to the background in this task. When making a sandwich or performing a similar task where the scene is stationary, almost all the fixations can be explained by the task needs (Hayhoe et al., 2003; Land

et al., 1999), but in the situation explored here, the world changes dynamically as the subject walks through it, the targets and obstacles are positioned randomly, and there is no simple step-by-step "recipe" for task completion. The experimental context is interesting for this reason as it seems to typify the challenges that the visual system normally handles effectively, and demands both top-down and bottom-up mechanisms. In this paper we have addressed only the top-down mechanisms, and those fixations that can be explained on the basis of the experimentally defined tasks. This leaves the remaining fixations on other parts of the room unexplained. It is likely that there is some intrinsic value to looking at objects in the room, regardless of instruction. Since this might be considered the job of vision in general, such fixations would be expected, and might be the basis of a general strategy of sampling the environment for potentially useful information. Alternatively, such fixations could be necessary for visual search; not all fixations land directly on objects of interest but could guide future fixations towards important information, perhaps reflecting strategies involving peripheral vision by looking at a location in between several objects instead of looking directly at one. Background fixations could also reflect strategies of getting global information within the room for navigation. It may be that the other fixations reflect additional "tasks" or modules that are relevant for navigation, and that we have not explicitly identified in our paradigm, such as evaluating global scene properties, obstacle density, etc. In any given situation we don't know what the specific modules might be, and that remains a challenge for future work. Alternatively, some of these other fixations might have a bottom-up origin, and understanding how bottom-up factors might play out in a context like this is a difficult problem (Tatler et al., 2011).

Conclusion

Understanding human sensorimotor behavior presents a unique challenge, largely because we do not know how behavior is organized. The Modular RL model of Sprague et al. (2007) attempts to make the problem more tractable by assuming that complex behavior can be composed by small sets of specialized component behaviors, or modules. This simplifies the problem to one of understanding which component behavior will be active at any particular time. Their model proposes that modules are chosen based on state uncertainty that could cause a reduction in expected reward, and that this updating process can be revealed by gaze target selection. In this paper, we attempt to test these core assumptions. First, we show that

manipulating both uncertainty and reward have a large effect on both walking behavior and eye movements. An increase in a module's relevance or uncertainty led to the predicted increase of eye moments to task relevant objects, supporting prior work on the topic. Irrelevant moving objects showed no significant increase of fixations despite being highly salient and relevant on other trials in the experiment. Uncertainty affected both the distance at which objects were fixated and the clearance subjects gave objects when planning their trajectories.

The results shed some light on the potential modularity of the combined set of tasks. Increasing the relevance or uncertainty of a module affected the behavior of that module as expected, but there were also some interactions between modules beyond what one would observe if they were completely independent. For instance, increasing the reward associated with targets also increased both looks to obstacles and obstacle avoidance. Teasing apart the causes and nature of this interaction necessitates further research. This approach of breaking down the complex problem of human behavior into independent modules thus provides valuable insights, but serves as a starting point for understanding, and perhaps as a strategy for exploring the added complexities of particular behaviors.

Keywords: active vision, eye movements, attention

References

Ballard, D., & Hayhoe, M. (2009). Modeling the role of task in the control of gaze. *Visual Cognition*, *17*(6/7), 1185–1204.

Bichot, N. P., & Schall, J. D. (1999). Saccade target selection in macaque during feature and conjunction visual search. *Vision Neuroscience*, *16*, 81–89.

Bisley, J. W., & Goldberg, M. E. (2010). Attention, intention, and priority in the parietal lobe. *Annual Review of Neuroscience*, *33*, 1–21.

Borji, A., & Itti, L. (2014). Computational models of attention: Top down and bottom up aspects. In A. C. Nobre & S. Kastner (Eds.), *The Oxford handbook of attention* (pp. 1122–1158). Oxford, UK: Oxford University Press.

Borji, A., Sihite, D. N., & Itti, L. (2011). Computational modeling of top-down visual attention in interactive environments. In J. Hoey, S. McKenna, & E. Trucco (Eds.), *Proceedings of the British Machine Vision Conference* (pp. 85.1–85.12). Durham, UK: BMVA Press.

Buswell, G. T. (1935). *How people look at pictures: A study of the psychology of perception in art.* Chicago, IL: University of Chicago Press.

Deubel, H., & Schneider, W. X. (1996). Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision Research, 36*(12), 1827–1837.

Droll, J., & Hayhoe, M. (2007). Trade-offs between gaze and working memory use. *Journal of Experimental Psychology: Human Perception and Performance, 33*(6), 1352–1365.

Droll, J., Hayhoe, M., Triesch, J., & Sullivan, B. (2005). Task demands control acquisition and maintenance of visual information. *Journal of Experimental Psychology: Human Perception and Performance, 31*(6), 1416–1438.

Eckstein, M. P., Schoonveld, W., Zhang, S., Mack, S. C., & Akbas, E. (2015). Optimal and human eye movements to clustered low value cues to increase decision rewards during search. *Vision Research, 113*(Pt B), 137–154, doi: 10.1016/j.visres.2015.05.016.

Faisal, A. A., & Wolpert, D. M. (2009). Near optimal combination of sensory and motor uncertainty in time during a naturalistic perception-action task. *Journal of Neurophysiology, 101*(4), 1901–1912.

Findlay, J. M., & Walker, R. (1999). A model of saccade generation based on parallel processing and competitive inhibition. *Behavioral and Brain Sciences, 22*, 661–721.

Foulsham, T., Tong, M., & Kingstone, A. (2016). The where, what and when of gaze allocation in the lab and the natural environment. *Vision Research, 51*(2011), 1920–1931.

Ghahghaei, S., & Verghese, P. (2015). Efficient saccade planning requires time and clear choices. *Vision Research, 113*, 125–136.

Gibson, B., Folk, C., Teeuwes, J., & Kingstone, A. (2008). Introduction to special issue on attentional capture. *Visual Cognition, 16*, 145–154.

Gottlieb, J. (2012). Attention, learning, and the value of information. *Neuron, 76*, 281–295.

Gottlieb, J., Hayhoe, M., Hikosaka, O., & Rangel, A. (2014). Attention, reward, and information-seeking. *Journal of Neuroscience, 34*(46), 15497–15504.

Hayhoe, M., & Ballard, D. (2014). Modeling task control eye movements. *Current Biology, 24*(13), 622–628.

Hayhoe, M. M., Shrivastava, A., Mruczek, R., & Pelz, J. B. (2003). Visual memory and motor planning in a natural task. *Journal of Vision, 3*(1):6, 49–63, doi: 10.1167/3.1.6. [PubMed] [Article]

Hitzel, E., Tong, M., Schutz, A., & Hayhoe, M. (2015). Objects in the peripheral visual field influence gaze location in natural vision. *Journal of Vision, 15*(12): 783, doi:10.1167/15.12.783. [Abstract]

Itti, L., &. Baldi, P. (2006). Modeling what attracts human gaze over dynamic natural scenes. In Harris L. & Jenkin M. (Eds.), *Computational vision in neural and machine systems* (pp. 183–200). Cambridge, MA: Cambridge University Press.

Jovancevic, J., & Hayhoe, M. (2009). Adaptive gaze control in natural environments. *Journal of Neuroscience, 29*(19), 6234–6238.

Jovancevic, J., Sullivan, B., & Hayhoe, M. (2006). Control of attention and gaze in complex environments. *Journal of Vision, 6*(12):9, 1431–1450, doi: 10.1167/6.12.9. [PubMed] [Article]

Kanan, C. M., Tong, M. H., Zhang, L., & Cottrell, G. W. (2009). SUN: Top-down saliency using natural statistics. *Visual Cognition, 17*, 979–1003.

Kit, D., Katz, L., Sullivan, B., Snyder, K., & Ballard, D. (2014). Eye movements, visual search and scene memory, in an immersive virtual environment. *PLoS ONE, 9*(4), e94362, doi:10.1371/journal.pone.0094362.

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature, 427*(6971), 244–247.

Kowler, E. (1990). (Ed.) Eye movements and their role in visual and cognitive processes, Vol. 4. In *Reviews of oculomotor research.* Amsterdam, The Netherlands: Elsevier.

Kowler, E., Anderson, E., Dosher, B., & Blaser, E. (1995). The role of attention in the programming of saccades. *Vision Research, 35*(13), 1897–1916.

Land, M. F., & Hayhoe, M. (2001). In what ways do eye movements contribute to everyday activities? *Vision Research, 41*, 3559–3566.

Land, M. F., Mennie, N., & Rusted, J. (1999). The roles of vision and eye movements in the control of

activities of daily living. *Perception*, 28(11), 1311–1328.

Land, M. F., & Tatler, B. W. (2009). *Looking and acting: Vision and eye movements in natural behaviour*. Oxford, UK: Oxford University Press.

Lee, D., Seo, H., & Jung, M. W. (2012). Neural basis of reinforcement learning and decision making. *Annual Review of Neuroscience*, 35, 287–308.

Li, C.-L., Aivar, M. P., Kit, D. M., Tong, M. H., & Hayhoe, M. M. (2016). Memory and visual search in naturalistic 2D and 3D environments. *Journal of Vision*, 16(8):9, 1–20, doi:10.1167/16.8.9. [PubMed] [Article]

Lin, J. Y., Franconeri, S., & Enns, J. T. (2008). Objects on a collision path with the observer demand attention. *Psychological Science*, 19, 686–692.

Maloney, L. T., & Zhang, H. (2010). Decision-theoretic models of visual perception and action. *Vision Research*, 50(23), 2362–2374.

Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, 434, 387–391.

Najemnik, J., & Geisler, W. S. (2008). Eye movement statistics in humans are consistent with an optimal search strategy. *Journal of Vision*, 8(3):4, 1–14, doi:10.1167/8.3.4. [PubMed] [Article]

Navalpakkam, V., & Itti, L. (2005). Modeling the influence of task on attention. *Vision Research*, 45, 205–231.

Navalpakkam, V., Koch, C., Rangel, A., & Perona, P. (2010). Optimal reward harvesting in complex perceptual environments. *Proceedings of the National Academy of Sciences, USA*, 107(11), 5232–5237.

Oliva, A., & Torralba, A. (2006). Building the gist of a scene: The role of global image features in recognition. *Progress in Brain Research*, 155, 23–36.

Pashler, H. (1998). *The psychology of attention*. Cambridge, MA: MIT Press.

Peters, R. J., & Itti, L. (2006). Computational mechanisms for gaze direction in interactive visual environments. In Proceedings of the ACM Symposium on Eye Tracking Research and Applications (pp. 27–32).

Reichle, E. D., Rayner, K., & Pollatsek, A. (2003). The EZ Reader model of eye-movement control in reading: Comparisons to other models. *Behavioral and Brain Sciences*, 26(04), 445–476.

Renninger, L. W., Verghese, P., & Coughlan, J. (2007). Where to look next? Eye movements reduce local uncertainty. *Journal of Vision*, 7(3):6, 1–17, doi:10.1167/7.3.6. [PubMed] [Article]

Rothkopf, C., Ballard, D., & Hayhoe, M. (2007). Task and context determine where you look. *Journal of Vision*, 7(14):16, 1–20, doi:10.1167/7.14.16. [PubMed] [Article]

Rothkopf, C., & Ballard, D. H. (2010), Module activation and credit assignment in reinforcement learning. *Frontiers in Psychology*, 1, 173.

Rothkopf, C., Hayhoe, M., & Ballard, D. (submitted). Control of sequential sensory motor decisions by subjective value in a natural task. *Current Biology*.

Schutz, A., Trommershauser, J., & Gegenfurtner, K. (2012). Dynamic integration of information about salience and value for saccadic eye movements. *Proceedings of the National Academy of Sciences, USA 109*, 7547–7552.

Schultz, W. (2000). Multiple reward signals in the brain, *Nature Reviews Neuroscience*, 1, 199–207.

Sprague, N., Ballard, D., & Robinson, A. (2007). Modeling embodied visual behaviors. *ACM Transactions on Applied Perception*, 4(2), 11.

Stritzke, M., Trommershauser, J., & Gegenfurtner, K. (2009). Effects of salience and reward information during saccadic decisions under risk. *Journal of the Optical Society of America A*, 25(11), B1–13.

Sullivan, B. T., Johnson, L. M., Rothkopf, C., Ballard, D., & Hayhoe, M. (2012). The role of uncertainty and reward on eye movements in a virtual driving task. *Journal of Vision 12*(13):19, 1–17, doi:10.1167/12.13.19. [PubMed] [Article]

Sutton, R., & Barto, A. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Tatler, B. W. (2007). The central fixation bias in scene viewing: Selecting an optimal viewing position independently of motor biases and image feature distributions. *Journal of Vision*, 7(14):4, 1–17, doi:10.1167/7.14.4. [PubMed] [Article]

Tatler, B., Hayhoe, M., Ballard, D. H., & Land, M. F. (2011). Eye guidance in natural vision: Reinterpreting salience. *Journal of Vision*, 11(5):5, 1–23, doi:10.1167/11.5.5. [PubMed] [Article]

't Hart, B. M., & Einhäuser, W. (2012). Mind the step: Complementary effects of an implicit task on eye and head movements in real-life gaze allocation. *Experimental Brain Research*, 223(2), 233–249.

Tong, M. H., & Hayhoe, M. (2014). Modeling uncertainty and intrinsic reward in a virtual walking task. *Journal of Vision*, 14(10): 5, doi:10.1167/14.10.5. [Abstract]

Torralba, A., Oliva, A., Castelhano, M., & Henderson, J. M. (2006). Contextual guidance of attention and eye movements in real world scenes: The role of

global features on object search. *Psychological Review*, *113*(4), 766–786.

Triesch, J., Ballard, D. H., Hayhoe, M. M., & Sullivan, B. T. (2003). What you see is what you need. *Journal of Vision*, *3*(1):9, 86–94, doi:10.1167/3.1.9. [PubMed] [Article]

Trommershäuser, J., Glimcher, P. W., & Gegenfurtner, K. (2009). Visual processing, learning and feedback in the primate eye movement system. *Trends in Neuroscience*, *32*, 583–590.

Verghese, P. (2012). Active search for multiple targets is inefficient. *Vision Research*, *74*, 61–71.

Zhang, H., Daw, N. D., & Maloney, L. T. (2015). Human representation of visuo-motor uncertainty as mixtures of orthogonal basis distributions. *Nature Neuroscience*, *18*(8), 1152–1158.

Zhang, L., Tong, M. H., Marks, T. K., Shan, H., & Cottrell, G. W. (2008). SUN: A Bayesian framework for saliency using natural statistics. *Journal of Vision*, *8*(7):32, 1–20, doi:10.1167/8.7.32. [PubMed] [Article]